arXiv:1101.3884v1 [quant-ph] 20 Jan 2011

Approximation algorithms for QMA-complete problems

Sevag Gharibian* Julia Kempe†

Abstract

Approximation algorithms for classical constraint satisfaction problems are one of the main research areas in theoretical computer science. Here we define a natural approximation version of the QMA-complete local Hamiltonian problem and initiate its study. We present two main results. The first shows that a non-trivial approximation ratio can be obtained in the class NP using product states. The second result (which builds on the first one), gives a polynomial time (classical) algorithm providing a similar approximation ratio for dense instances of the problem. The latter result is based on an adaptation of the "exhaustive sampling method" by Arora et al. [AKK99] to the quantum setting, and might be of independent interest.

1 Introduction and Results

In the last few years, the quantum analog of the class NP, the class QMA [KSV02], has been extensively studied, and several QMA-complete problems have been found [Liu06, Bra06, LCV07, BS07, Ros09, JGL10, SV09, WMN10]. Arguably the most important QMA-complete problem is the k-local Hamiltonian problem [KSV02, KR03, OT08, KKR06, AGIK09]. Here, the input is a set of Hamiltonians (Hermitian matrices), each acting on at most k-qubits each. The task is to determine the largest eigenvalue of the sum of these Hamiltonians. This problem generalizes the central NP-hard problem MAX-k-CSP, where we are given a set of Boolean constraints on k variables each, with the goal to satisfy as many constraints as possible. The local Hamiltonian problem is of significant interest to complexity theorists and to physicists studying properties of physical systems alike (e.g. [BV05, AvDK$^+$07, BDOT08, AALV09, CV09, LLM$^+$10, SC10]).

Moving to the classical scenario, the theory of NP-completeness is one of the great success stories of classical computational complexity [AB09]. It was soon realized that many natural optimization problems are NP-hard, and are hence unlikely to have polynomial time algorithms. A natural question (both in theory and in practice) is to look for polynomial time algorithms that produce solutions that are close to optimum. More precisely, one says that an algorithm achieves an *approximation ratio* of $c \in [0,1]$ for a certain maximization problem if on all inputs, the value of the algorithm's output is at least c times that of the optimum solution. The closer c is to 1, the better the approximation. The investigation of approximation algorithms is, after decades of heavy research, still a very active area (e.g., [Hoc97, Vaz01]). For many central NP-hard problems, tight polynomial time approximation algorithms are known.

In the context of QMA-complete problems, it is thus natural to search for approximation algorithms for these problems, and in particular for the local Hamiltonian problem. The question we address here is: *How well can one efficiently approximate the k-local Hamiltonian problem?*

---

[*] David R. Cheriton School of Computer Science and Institute for Quantum Computing, University of Waterloo, Waterloo N2L 3G1, Canada. Supported by the Natural Sciences and Engineering Research Council of Canada (NSERC), NSERC Michael Smith Foreign Study Supplement, and by EU-Canada Transatlantic Exchange Partnership programme.

[†] Blavatnik School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel and CNRS & LRI, University of Paris-Sud, Orsay, France. Supported by an Individual Research Grant of the Israeli Science Foundation, by European Research Council (ERC) Starting Grant QUCO and by the Wolfson Family Charitable Trust.

It should be noted that a large host of heuristics has been developed in the physics community to approximate properties of local Hamiltonian systems (see, e.g., [CV09] for a survey) and this area is extremely important in the study of physical systems. However, the systematic complexity theoretic study of approximation algorithms for QMA-complete problems is still very much in its infancy, and our work is one of the first steps in this research direction. We note that there has been a lot of interest in recent years [AALV09, Aar06] in establishing a so-called quantum PCP theorem (e.g. [AS98, ALM$^+$98]), which amounts to showing that for some constant $c < 1$ close enough to 1, approximating the k-local Hamiltonian (or related problems) to within c is QMA-hard. Our results can also be seen as a natural continuation of that investigation.

Our results: Let us start by precisely defining the optimization version of the local Hamiltonian problem, which is parameterized by two integers k and d, which we always think of as constants.

Definition 1 (MAX-k-local Hamiltonian problem on d-level systems (qudits))**.** *An instance of the problem consists of a collection of $\binom{n}{k}$ Hermitian matrices, one for each subset of k qudits. The matrix $H_{i_1,\dots,i_k}$ corresponding to some $1 \le i_1 \le \cdots \le i_k \le n$ is assumed to act on those qudits*[1]*, to be positive semidefinite, and to have operator norm at most* 1*. We call any pure or mixed state ρ on n qudits an* assignment *and define its* value *to be $TrH\rho$ where $H = \sum_{i_1,\dots,i_k} H_{i_1,\dots,i_k}$. The goal is to find the largest eigenvalue of H (denoted* OPT*), or equivalently, the maximum value obtained by an assignment. We say that an algorithm provides an* approximation ratio *of $c \in [0,1]$ if for all instances, it outputs a value that is between $c \cdot \mathrm{OPT}$ and* OPT.

This definition, we believe, is the natural quantum analog of the MAX-k-CSP problem. We note that it differs slightly from the usual definition of the k-local Hamiltonian problem. Namely, we consider maximization (as opposed to minimization), and also restrict the terms of H to be positive semidefinite, and have norm at most 1. As long as one considers the *exact* problem, these assumptions are without loss of generality, and do not affect the definition, as seen by simply scaling the Hamiltonians and adding multiples of identity as necessary. However, when dealing with the *approximation* version, these assumptions are important for the problem to make sense; for instance, one cannot meaningfully talk about approximation ratios if the optimum can take both negative and positive values. That is why we require the terms to be positive semidefinite. The requirement that the terms have operator norm at most 1 does not affect the problem and later allows us to conveniently define dense instances. Finally, changing the maximization to a minimization would lead to an entirely different approximation problem: the quantum analogue of MIN-CSP (e.g. [KSTW01]). Minimization problems are, generally speaking, harder than maximization problems, and we leave this research direction for future work.

Before stating our results, we remark that there is a trivial way to get a d^{-k}-approximation for MAX-k-local Hamiltonian. Observe that the maximally mixed state has at least d^{-k} overlap with the reduced density matrix of the optimal assignment on any k particles. A similar thing holds classically, where a random assignment gives (in expectation) a d^{-k} approximation of MAX-k-CSP. We now describe our two main results.

Approximation by product states. One inherently quantum property of the local Hamiltonian problem is the fact that the optimal state might in general be highly entangled (and hence not efficiently describable in polynomial time). This is why we do not require outputting the assignment itself in the above definition. If, however, the optimal assignment (or some other good assignment) was guaranteed to be a *tensor product state*, then we could describe it efficiently. The following theorem shows just that.

Theorem 2. *For an instance of MAX-k-local Hamiltonian with optimal value* OPT*, there is a (pure) product state assignment that has value at least* OPT/d^{k-1}.

---

[1] Terms acting on less than k qudits can be incorporated by tensoring them with the identity.

This result is *tight* for product states in the case of 2-local Hamiltonians (we remark that 2-local Hamiltonians are often the most relevant case from a physics perspective). For example, consider the Hamiltonian on 2-qubits that projects onto the EPR state $\frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. It is easy to see that no product state achieves value more than $1/2$. For general k, we can only show that product states cannot achieve an approximation ratio greater than $1/d^{\lfloor k/2 \rfloor}$ (see Sec. 3).

If we could efficiently find the best product state assignment, we would obtain an algorithm achieving a non-trivial d^{-k+1} approximation ratio. Unfortunately, this problem is NP-complete, since it would allow one to solve (e.g.) the special case of MAX-k-SAT, and hence we do not have such an algorithm. Still, the theorem has the following interesting implication: It shows that unless $\mathrm{NP} = \mathrm{QMA}$, approximating the local Hamiltonian problem to within a factor less than d^{-k+1} is not QMA-hard. This follows simply because product states have polynomial size classical descriptions.

A polynomial time approximation algorithm for dense instances. Our second result gives a classical polynomial time approximation algorithm for *dense* instances of the local Hamiltonian problem. This result is perhaps our technically most challenging one, and we hope the techniques we develop might turn out useful elsewhere.

Dense instances of classical constraint satisfaction problems have been studied in depth (see e.g. [dlV96, FK96, GGR98, AKK99, dlVK00, AdlVKK02, BdlVK03, dlVKKV05]). Our result is inspired by work of Arora et al. [AKK99] who provide a polynomial time approximation scheme, or PTAS (i.e., an efficient $1 - \varepsilon$ approximation algorithm for any fixed $\varepsilon > 0$), for several types of dense constraint satisfaction problems. In the classical case, dense (for 2-local constraints) simply means that the average degree in the constraint graph is $\Omega(n)$, or equivalently, that the optimum is $\Omega(n^2)$. In analogy, we define an instance of MAX-k-local Hamiltonian to be *dense* if $\mathrm{OPT} = \Omega(n^k)$, or equivalently, if $Tr(H\frac{Id}{d^n}) = \Omega(n^k)$.[2]

It is not hard to see that the (exact) dense local Hamiltonian problem remains QMA-hard (see Sec. 3.3). The dense case might be of practical interest to physicists who study systems of particles by incorporating all possible interactions between them. Note that such instances are dense even if the interactions between particles are weak, so long as the interaction strengths are constants independent of n. Our second main result is the following:

Theorem 3. *For all $\varepsilon > 0$ there is a polynomial time $(1/d^{k-1} - \varepsilon)$ approximation algorithm for the dense MAX-k-local Hamiltonian problem over qudits.*

Thm. 3 follows immediately by combining Thm. 2 with the following theorem, which gives an approximation scheme for the problem of optimizing over the set of product states.

Theorem 4. *For all $\varepsilon > 0$ there is a polynomial time algorithm for dense MAX-k-local Hamiltonian that outputs a product state assignment with value within $1 - \varepsilon$ of the value of the best product state assignment.*

Proof ideas and new tools: The proofs of Thm. 2 and Thm. 4 are independent and employ different techniques. To show the product state approximation guarantee, we show a slightly stronger statement: For *any* assignment $|\Psi\rangle$, there is a way to construct a product assignment of at least d^{-k+1} its value. The proof is constructive (given $|\Psi\rangle$): we use a type of recursive Schmidt decomposition of $|\Psi\rangle$ to obtain a mixture of product states whose value is off by at most the desired approximation factor (see Sec. 2).

Our second result is technically more challenging and introduces a few new ideas to this problem, inspired by work of Arora et al. [AKK99] in the classical setting. We illustrate the main ideas for MAX-2-local Hamiltonian on n qubits. Recall that our goal is to find a PTAS for the local Hamiltonian problem

---

[2] The equivalence follows from the fact that the mixed state assignment Id/d^n has value between OPT and OPT/d^k.

over product states. The value of the optimal product state assignment, OPT_P, can be written

$$\mathrm{OPT}_P \quad = \quad \max \quad \sum_{i=1}^{n} \sum_{j \in N(i)} Tr(H_{i,j}(\rho_i \otimes \rho_j)) \quad \text{s.t.} \quad \rho_i \succeq 0 \text{ and } \mathrm{Tr}(\rho_i) = 1 \quad \text{for } 1 \leq i \leq n, \tag{1}$$

where $N(i)$ is the set of indices j for which a local Hamiltonian term $H_{i,j}$ is present. We might call this a *quadratic* semidefinite program, as the maximization is quadratic in the ρ_i (and as such not efficiently solvable). Note, however, that if the terms in the maximization were *linear*, then we would obtain a semidefinite program (SDP), which is efficiently solvable [GLS93]. To "linearize" our optimization, we use the "exhaustive sampling method" developed by Arora et al. [AKK99] (a method which was later key in many developments in property testing, e.g. [GGR98]). We write each Hamiltonian term in a basis that separates its two qubits, for instance the Pauli basis $\{\sigma_0, \sigma_1, \sigma_2, \sigma_3\}$, $H_{i,j} = \sum_{k,l=0}^{3} \alpha_{kl}^{ij} \sigma_k \otimes \sigma_l$. For $i = 1, \ldots, n$ and $k = 0, 1, 2, 3$, define

$$c_k^i := \sum_{j \in N(i)} \sum_{l} \alpha_{kl}^{ij} Tr(\sigma_l \rho_j).$$

If we knew the values of c_k^i for the optimal ρ_i, then solving the SDP below would yield the optimal ρ_i:

$$\begin{aligned} \max \quad \sum_{i=1}^{n} \sum_{k=0}^{3} c_k^i Tr(\sigma_k \rho_i) \quad \text{s.t.} \quad & \rho_i \succeq 0 \text{ and } \mathrm{Tr}(\rho_i) = 1 && \text{for } 1 \leq i \leq n, \\ & \sum_{j \in N(i)} \sum_{l} \alpha_{kl}^{ij} Tr(\sigma_l \rho_j) = c_k^i && \text{for } 1 \leq i \leq n \text{ and } 0 \leq k \leq 3. \end{aligned} \tag{2}$$

Of course, this reasoning is circular, as in order to obtain the c_k^i we need the optimal ρ_i. The crucial idea is now to use *sampling* to *estimate* the c_k^i. More precisely, assume for a second that we could sample $O(\log n)$ of the ρ_i randomly from the optimal assignment. Then, by standard sampling bounds, with high probability over the choice of the sampled qubits we can estimate all the c_k^i to within an additive error $\pm \varepsilon n$ for some ε. If we had these estimates a_k^i for the c_k^i, we could solve the SDP above with the slight modification that the last constraint should be $a_k^i - \varepsilon n \leq \sum_{j \in N(i)} \sum_l \alpha_{kl}^{ij} Tr(\sigma_l \rho_j) \leq a_k^i + \varepsilon n$. With high probability over the sampled qubits, this SDP will give a solution that is within an additive εn^2 of the optimal one (more subtle technicalities and all calculations can be found in Sec. 3). Moreover, it is possible to derandomize the sampling procedure to obtain a deterministic algorithm (Sec. 3.3).

Of course, we are still in the realm of wishful thinking, because in order to sample from the optimal solution, we would need to know it, which is precisely what we set out to do. However, the number of qubits we wish to sample is only *logarithmic* in the input size. Thus, to simulate the sampling procedure, we can pick a random subset of $O(\log n)$ qubits, and simply *iterate* through all possible assignments on them (with an appropriate δ-net over the density matrices, which incurs a small additional error) in polynomial time! Our algorithm then runs the SDP for each iteration, and we are guaranteed that at least one iteration will return a solution within εn^2 of the optimal one. Because the denseness assumption guarantees that OPT_P is $\Omega(n^2)$, our additive approximation turns into a factor $(1 - \varepsilon)$-approximation, as desired. All details, the runtime of the algorithm and error bounds for the general k-local case on qudits are given in Sec. 3.

Previous and related work: We note that many heuristics have been developed in the physics community to approximate properties of local Hamiltonian systems and this area is extremely important in the study of physical systems (e.g. [Whi92, Whi93, ÖR95, RÖ97, Sch05, PWKE98, CV09]). Our focus here is, however, on rigorous bounds on the approximation guarantee of algorithms for the general problem. In this area, to our knowledge, few results are known. In a first result on polynomial time approximation algorithms, Bansal, Bravyi and Terhal [BBT09] give a PTAS for a special class of the local Hamiltonian problem, so called quantum Ising spin glasses, for the case where the instance is on a planar graph (and in particular of bounded degree). This PTAS is obtained by dividing the graph into constant size chunks, which can be solved directly, and ignoring the constraints between chunks (this incurs an error proportional to the number of such constraints, which is small because the graph is planar). More recently, there has been work proving rigorous approximations to ground states of one-dimensional quantum systems under well-defined conditions using techniques such as density matrix renormalization group [AAI10, SC10]. To our knowledge, we are the first to establish a bound on the approximation factor by optimizing over the set of product states.

Discussion and open questions: Our two results give approximations to the local Hamiltonian problem. Although at first glance, our approximation ratio of $1/d^{k-1}$ may appear an incremental improvement over the trivial random assignment strategy, there are two important notes that should be kept in mind: The first is that many classical NP-hard problems, such as MAX-3-SAT (a special case of MAX-k-CSP where each constraint is the disjunction ("OR") of k variables or their negation), are *approximation resistant* (e.g. [Hås07, AM08]), meaning that unless P$=$NP, there do not even exist non-trivial approximation ratios beyond the random assignment strategy. For example, for MAX-3-SAT it is NP-hard to do better than the approximation ratio of $7/8$ achieved by random assignment [Hås97]. Thus, showing the existence of a non-trivial approximation ratio is typically a big step in the classical setting. Moreover, it could have been conceivable that for MAX-k-local Hamiltonian, analogously to MAX-3-SAT, outperforming the random assignment strategy would have been *QMA-hard*. Yet our results show that unless NP$=$QMA, this is not the case. The second important note that should be kept in mind is that the currently *best* approximation algorithm for MAX-k-CSP gives an approximation ratio of only about $0.44k/2^k$ for $k > 2$ [CMM07] (for $k = 2$, one can achieve 0.874 [LLZ02]) and this is, moreover, essentially the best possible under a plausible complexity theoretic conjecture (namely, the unique games conjecture [Kho02]) [Tre98, Has05, ST06, AM08]. This is to be contrasted with our $2/2^k$-approximation ratio for the case of $d = 2$ (i.e. qubit systems), which we show can be achieved by product state assignments for *arbitrary* (i.e. even non-dense) MAX-k-local Hamiltonian instances. This raises the important open question: is our approximation ratio tight?

Our product state approximation shows that approximating the local Hamiltonian problem to within d^{-k+1} is in NP. It would be interesting to know if this approximation ratio could also be achieved in polynomial time. If not, it might lead to an intriguing state of affairs where for low approximation ratios the problem is efficiently solvable, for medium ratios it is in NP but not efficiently solvable, and for high ratios it is QMA-hard (assuming a quantum PCP theorem exists).

To obtain our results for the case of dense local Hamiltonians, we have introduced the exhaustive sampling technique of Arora et al. [AKK99] to the setting of low-degree semidefinite programs. We linearize such programs using exhaustive sampling in combination with a careful analysis of the error coming from working with δ-nets on density matrices. We remark that it seems we cannot simply apply the results of [AKK99] for *smooth Polynomial Integer Programs* as a black-box to our setting. This is due to our aforementioned need for a δ-net, as well as the requirement that our assignment be a positive semidefinite operator. We address the latter issue by extending the techniques of [AKK99] to the realm of positive semidefinite programs by introducing the notion of "degree-k inner products" over Hermitian operators to generalize the concept of degree-k polynomials over real numbers, and performing the more complex analysis that ensues. We hope that this technique will be of much wider applicability, particularly considering the growing use of semidefinite programs in numerous areas of quantum computing and information (e.g. [DPS04, JJUW10, LMRS10]).

Another open question is whether similar ideas can be used to approximate other QMA-complete problems, such as the consistency problem [Liu06]. Moreover, can we obtain polynomial time algorithms without the denseness assumption? And are there special cases of the local Hamiltonian problem for which there is a PTAS (other than for planar Ising spin glasses [BBT09])? Of course, we do not expect a PTAS for all instances of the local Hamiltonian problem, as this would contradict known hardness results for special classical cases of the problem. However, perhaps there exist other classes of physically relevant instances of the problem for which a PTAS does exist.

Structure of this paper: In Sec. 2, we prove our result on product state approximations (Thm. 9 and the ensuing proof of Thm. 2), show its tightness in the 2-local case and provide the upper bound of $d^{-\lfloor k/2 \rfloor}$ for the best possible approximation by product states. Sec. 3 gives our polynomial time approximation algorithm and develops the general sampling and SDP-based technique we use. It also shows that the dense local Hamiltonian problem remains QMA-complete.

Notation: We use $A \succeq 0$ to say operator A is positive semidefinite, and denote by $L(\mathcal{X})$, $H(\mathcal{X})$, and $D(\mathcal{X})$ the sets of linear, Hermitian, and density operators acting on complex Euclidean space $\mathcal{X}$, respectively. We denote the Frobenius and operator norms of $A \in L(\mathcal{X})$ as $\| A \|_{\mathrm{F}} = \sqrt{\mathrm{Tr}(A^\dagger A)}$ and $\| A \|_\infty = \max_{|x\rangle \in \mathcal{X} \text{ s.t. } \| x \|_2 = 1} \| A|x\rangle \|_2$, respectively.

2 Product states yield a $1/d^{k-1}$-approximation for qudits

We now show that product state assignments achieve a non-trivial approximation ratio for MAX-*k*-local Hamiltonian , i.e. Thm 2. The heart of our approach is what we call the *Mixing Lemma* (Lem. 7), which we use to prove Thm. 9. Thm. 2 will then easily follow. At the end of the section, we discuss the tightness of the approximation guarantee given by Thm. 2. We begin with two definitions.

Definition 5 (Recursive Schmidt Decomposition (RSD))**.** We define the *recursive Schmidt decomposition* of a state $|\psi\rangle \in (\mathbb{C}^d)^{\otimes n}$ as the expression obtained by recursively applying the Schmidt decomposition on each qudit from 1 to $n-1$ inclusive[3]. For example, the RSD for 3-qubit $|\psi\rangle$ is

$$|\psi\rangle = \alpha_1 |a_1\rangle \otimes (\beta_1 |b_1\rangle |c_1\rangle + \beta_2 |b_2\rangle |c_2\rangle) + \alpha_2 |a_2\rangle \otimes (\beta'_1 |b'_1\rangle |c'_1\rangle + \beta'_2 |b'_2\rangle |c'_2\rangle),$$

for $\alpha_1^2 + \alpha_2^2 = \beta_1^2 + \beta_2^2 = {\beta'_1}^2 + {\beta'}_2^2 = 1$, $\{|a_i\rangle\}_i$ an orthonormal basis for qubit 1, $\{|b_i\rangle\}_i$ and $\{|b'_i\rangle\}_i$ orthonormal bases for qubit 2, and $\{|c_i\rangle\}_i$ and $\{|c'_i\rangle\}_i$ orthonormal bases for qubit 3.

Definition 6 (Schmidt cut)**.** For any $|\psi\rangle \in (\mathbb{C}^d)^{\otimes n}$ with Schmidt decomposition $|\psi\rangle = \sum_{i=1}^d \alpha_i |w_i\rangle |v_i\rangle$, where $|w_i\rangle \in \mathbb{C}^d$ and $|v_i\rangle \in (\mathbb{C}^d)^{\otimes n-1}$, and for any $|\phi\rangle \in (\mathbb{C}^d)^{\otimes m}$, we refer to the expansion $|\phi\rangle \otimes \left(\sum_{i=1}^d \alpha_i |w_i\rangle |v_i\rangle\right)$ as the *Schmidt cut* at qudit $m+1$. We say that a projector Π *crosses* this Schmidt cut if Π acts on qudit $m+1$ and at least one qudit $i \in \{m+2, \ldots, m+n\}$.

The heart of our approach is the following Mixing Lemma, which provides, for *any* assignment $|\psi\rangle \in (\mathbb{C}^d)^{\otimes n}$, an explicit construction through which the entanglement across the first Schmidt cut of $|\psi\rangle$ can be eliminated, while maintaining at least a $(1/d)$-approximation ratio relative to the value $|\psi\rangle$ achieves against any local Hamiltonian $H \in H((\mathbb{C}^d)^{\otimes n})$.

Lemma 7 (Mixing Lemma)**.** *Given state $|\psi\rangle$ on n qudits with Schmidt cut on qudit* 1 *given by* $|\psi\rangle = \sum_{i=1}^d \alpha_i |w_i\rangle |v_i\rangle$*, where* $|w_i\rangle \in \mathbb{C}^d$ *and* $|v_i\rangle \in (\mathbb{C}^d)^{\otimes n-1}$*, define* $\rho := \sum_{i=1}^d \alpha_i^2 |w_i\rangle\langle w_i| \otimes |v_i\rangle\langle v_i|$*. Then, given projector* Π *acting on some subset* $\mathcal{S}$ *of the qudits, if* Π *crosses the Schmidt cut, then* $\mathrm{Tr}(\Pi\rho) \geq \frac{1}{d}\mathrm{Tr}(\Pi|\psi\rangle\langle\psi|)$*. Otherwise,* $\mathrm{Tr}(\Pi\rho) = \mathrm{Tr}(\Pi|\psi\rangle\langle\psi|)$.

---

[3] This definition is relative to some fixed ordering of the qudits. The specific choice of ordering is unimportant in our scenario, as any decomposition output by such a process suffices to prove Thm. 2.

Proof. Case 2 follows easily by noting that the given Schmidt decomposition of $|\psi\rangle$ implies $\mathrm{Tr}_1(\rho) = \mathrm{Tr}_1(|\psi\rangle\langle\psi|)$ and $\mathrm{Tr}_{2,\dots,n}(\rho) = \mathrm{Tr}_{2,\dots,n}(|\psi\rangle\langle\psi|)$. To prove case 1, we observe by straightforward expansion that

$$\mathrm{Tr}(\Pi|\psi\rangle\langle\psi|) = \mathrm{Tr}(\Pi\rho) + \sum_{i<j} \alpha_i\alpha_j \langle w_i|\langle v_i|\Pi|w_j\rangle|v_j\rangle + \alpha_i\alpha_j\langle w_j|\langle v_j|\Pi|w_i\rangle|v_i\rangle. \tag{3}$$

Then, by defining for each i vector $|a_i\rangle := \alpha_i\Pi|w_i\rangle|v_i\rangle$, we have that

$$\sum_{i<j} \alpha_i\alpha_j \langle w_i|\langle v_i|\Pi|w_j\rangle|v_j\rangle + \alpha_i\alpha_j\langle w_j|\langle v_j|\Pi|w_i\rangle|v_i\rangle = \sum_{i<j}\langle a_i|a_j\rangle + \langle a_j|a_i\rangle.$$

Applying the facts that $\Pi^2 = \Pi$ and $\langle a|b\rangle + \langle b|a\rangle \leq \| \, |a\rangle \, \|_2^2 + \| \, |b\rangle \, \|_2^2$ for $|a\rangle, |b\rangle \in (\mathbb{C}^d)^{\otimes n}$ implies

$$\sum_{i<j}\langle a_i|a_j\rangle + \langle a_j|a_i\rangle \leq \sum_{i<j} \| \, |a_i\rangle \, \|_2^2 + \| \, |a_j\rangle \, \|_2^2 = (d-1)\sum_i \alpha_i^2 \langle w_i|\langle v_i|\Pi|w_i\rangle|v_i\rangle = (d-1)\mathrm{Tr}(\Pi\rho),$$

from which the claim follows. ☐

The following simple extension of Lem. 7 simplifies our proof of Thm. 9.

Corollary 8. *Define $|\psi'\rangle := |\phi\rangle \otimes |\psi\rangle$, where $|\phi\rangle \in (\mathbb{C}^d)^{\otimes m}$ for $m > 0$ and $|\psi\rangle$ is defined as in Lem. 7, and let $\rho \in D(\mathbb{C}^d)^{\otimes n}$ be obtained from $|\psi\rangle$ as in Lem. 7. Then, for any projector Π acting on a subset $\mathcal{S}$ of the qudits, if Π crosses the Schmidt cut of $|\psi'\rangle$ at qudit $m+1$, we have $\mathrm{Tr}(\Pi|\phi\rangle\langle\phi| \otimes \rho) \geq \frac{1}{d}\mathrm{Tr}(\Pi|\psi'\rangle\langle\psi'|)$. Otherwise, $\mathrm{Tr}(\Pi|\phi\rangle\langle\phi| \otimes \rho) = \mathrm{Tr}(\Pi|\psi'\rangle\langle\psi'|)$.*

Proof. Immediate by applying the proof of Lem. 7 with the following modifications: (1) Define $|a_i\rangle := \alpha_i\Pi|\phi\rangle|w_i\rangle|v_i\rangle$, and (2) if $\mathcal{S} \subseteq \{1,\dots,m\} \cup \{m+2,\dots,m+n\}$ (i.e. this is one of two ways for Π not to cross the cut — the other way is for $\mathcal{S} \subseteq \{1,\dots,m+1\}$), observe that by the same arguments as in Lem. 7 for case 2 and the product structure between $|\phi\rangle$ and $|\psi\rangle$ in $|\psi'\rangle$ that $\mathrm{Tr}_{m+1}(|\phi\rangle\langle\phi| \otimes \rho) = \mathrm{Tr}_{m+1}(|\psi'\rangle\langle\psi'|)$. ☐

Lemma 7 shows that the state ρ obtained by *mixing* the d Schmidt vectors of $|\psi\rangle$, as opposed to taking their *superposition*, suffices to achieve a $(1/d)$-approximation across the first Schmidt cut. By iterating this argument over *all* $n-1$ Schmidt cuts, we now prove that a mixture of all (product) states appearing in the RSD of $|\psi\rangle$ achieves an approximation ratio of $1/d^{k-1}$.

Theorem 9. *For any n-qudit assignment $|\psi\rangle$ with RSD $|\psi\rangle = \sum_{i=1}^{d^{n-1}} \sqrt{p_i}|\phi_i\rangle$, where $\sum_i p_i = 1$ and $\{|\phi_i\rangle\}_{i=1}^{d^{n-1}}$ is a set of orthonormal product vectors in $(\mathbb{C}^d)^{\otimes n}$, define $\rho := \sum_{i=1}^{d^{n-1}} p_i|\phi_i\rangle\langle\phi_i|$. Then, for any projector Π acting on some subset $\mathcal{S} \subseteq \{1,\dots,n\}$ of qudits with $|\mathcal{S}| = k$, we have $\mathrm{Tr}(\Pi\rho) \geq \frac{1}{d^{k-1}}\mathrm{Tr}(\Pi|\psi\rangle\langle\psi|)$.*

Proof. Let Π be a projector with $|\mathcal{S}| = k$, and define $\mathbf{c} \in \{0,1\}^{n-1}$ such that $\mathbf{c}(j) = 1$ iff Π crosses the Schmidt cut at qudit j. For example, if Π acts on qudits $\{1,2\}$, then $\mathbf{c} = (1,0,\dots,0)$. Note that in general $\| \mathbf{c} \|_1 = k-1$. We proceed by iteratively stepping through each Schmidt cut in the RSD of $|\psi\rangle$. Let $\rho^{(0)} := |\psi\rangle\langle\psi|$, and consider first the cut at qudit 1, i.e. $|\psi\rangle = \sum_{i=1}^d \alpha_i|w_i\rangle|v_i\rangle$, for $|w_i\rangle \in \mathbb{C}^d$ and $|v_i\rangle \in (\mathbb{C}^d)^{\otimes n-1}$. Defining $\rho^{(1)} := \sum_{i=1}^d \alpha_i^2|w_i\rangle\langle w_i| \otimes |v_i\rangle\langle v_i|$, we have by Lem. 7 that

$$\mathrm{Tr}(\Pi|\psi\rangle\langle\psi|) \leq d^{\mathbf{c}(1)}\mathrm{Tr}(\Pi\rho^{(1)}), \tag{4}$$

i.e. we lose a factor of $1/d$ iff Π crosses the first cut.

Moving on to the second Schmidt cut, consider the state $|w_1\rangle|v_1\rangle \in \mathbb{C}^d \otimes (\mathbb{C}^d)^{\otimes n-1}$ appearing in the expression for $\rho^{(1)}$. Observe that it satisfies the preconditions for Cor. 8 with $m = 1$. Hence, via Cor. 8 there exists a state σ_1 acting on qudits $\{2,\dots,n\}$ such that $\mathrm{Tr}(\Pi|w_1\rangle\langle w_1|\otimes|v_1\rangle\langle v_1|)\leq d^{\mathbf{c}(2)}\mathrm{Tr}(\Pi|w_1\rangle\langle w_1|\otimes\sigma_1)$. We can analogously find states σ_i corresponding to $|w_i\rangle|v_i\rangle$ for all $1\leq i\leq d$. Thus,

$$\mathrm{Tr}(\Pi\rho^{(1)})=\sum_{i=1}^{d}\alpha_i^2\mathrm{Tr}(\Pi|w_i\rangle\langle w_i|\otimes|v_i\rangle\langle v_i|)\leq d^{\mathbf{c}(2)}\left[\sum_{i=1}^{d}\alpha_i^2\mathrm{Tr}(\Pi|w_i\rangle\langle w_i|\otimes\sigma_i)\right]. \tag{5}$$

Hence, by defining $\rho^{(2)}:=\sum_{i=1}^{d}\alpha_i^2|w_i\rangle\langle w_i|\otimes\sigma_i$, we have via Eqns. (4) and (5) that

$$\mathrm{Tr}(\Pi|\psi\rangle\langle\psi|)\leq d^{\mathbf{c}(1)+\mathbf{c}(2)}\mathrm{Tr}(\Pi\rho^{(2)}).$$

Since by Cor. 8, the σ_i are mixtures of Schmidt vectors from the second Schmidt cut, we can now iteratively apply the same procedure to the (at most d^2) pure states appearing in the expression for $\rho^{(2)}$ when considering the third Schmidt cut. Note in particular that each of these terms will have a product structure between qudits $\{1,2\}$ and $\{3,\dots,n\}$, as required by Cor. 8 for the next iteration.

More generally, when considering the pth Schmidt cut, we apply Cor. 8 with $m=p-1$ to each of the at most d^{p-1} terms appearing in the expansion of $\rho^{(p-1)}$. We continue iterating in this fashion until we have exhausted all $n-1$ Schmidt cuts, at which point the resulting mixture $\rho^{(n-1)}$ we are left with is in fact the ρ from the statement of the claim (seen by noting that our procedure essentially iteratively computes the RSD of $|\psi\rangle$, mixing the Schmidt vectors it computes at each step). Moreover, due to the repeated application of Cor. 8, we have

$$\mathrm{Tr}(\Pi|\psi\rangle\langle\psi|)\leq d^{\|\mathbf{c}\|_1}\mathrm{Tr}(\Pi\rho^{(n-1)}). \tag{6}$$

Recalling that $\|\mathbf{c}\|_1=k-1$ completes the proof. ☐

Proof of Thm. 2: Simply apply Thm. 9 to each projector in the spectral decompositions of each (positive semidefinite) H_i in our MAX-k-local Hamiltonian instance $H=\sum_i H_i$, and let $|\psi\rangle$ denote the optimal assignment for H. It is important to note that we can exploit Thm. 9 in this fashion due to the fact that the ρ constructed by Thm. 9 is *independent* of the projector Π — i.e. for any fixed $|\psi\rangle$ and k, the state ρ provides the same approximation ratio against *any* k-local projector Π encountered in the spectral decompositions of the H_i. Finally, note that one can find a *pure* product state achieving this approximation guarantee since ρ is a convex mixture of pure product states.

Upper bound of $d^{-\lfloor\frac{k}{2}\rfloor}$ for product state approximations. Is the result of Thm. 2 tight? In the case of MAX-2-local Hamiltonian on qudits, yes — consider a single clause projecting onto the maximally entangled state $\frac{1}{\sqrt{d}}\sum_i|ii\rangle$, for which a product state achieves value at most $1/d$. On the other hand, for MAX-3-local Hamiltonian on qubits, the worst case clause for a 3-qubit product state assignment is the projector onto the state $|W\rangle=\frac{1}{\sqrt{3}}(|001\rangle+|010\rangle+|100\rangle)$ [TWP09]. But here product states achieve value $4/9$ [WG03], implying the bound of $1/4$ from Thm. 2 is not tight.

A simple construction shows that the true optimal ratio is upper bounded by $d^{-\lfloor\frac{k}{2}\rfloor}$. To see this, consider a single clause which is the tensor product of maximally entangled bipartite states[4]. For example, for $n=4$, consider the clause $|\phi^+\rangle\langle\phi^+|\otimes|\phi^+\rangle\langle\phi^+|$, where $|\phi^+\rangle=\frac{1}{\sqrt{2}}(|00\rangle+|11\rangle)$. The maximum value a product state can attain is $1/4$, as claimed. In the qubit setting ($d=2$), one can further improve this construction for odd k by replacing the term $|\phi^+\rangle\langle\phi^+|\otimes I$ on the last three qubits with $|W\rangle\langle W|$. For example, for $k=5$, setting our instance to be the clause $|\phi^+\rangle\langle\phi^+|\otimes|W\rangle\langle W|$ yields an upper bound of $(1/2)(4/9)=2/9<1/4=d^{-\lfloor\frac{k}{2}\rfloor}$ (where we again use the value $4/9$ for $|W\rangle$ from the previous paragraph). For general odd $k>1$, this improved bound generalizes to $2^{\frac{-k+7}{2}}/9$.

---

[4] For odd k, we assume the odd qudit out projects onto the identity.

3 Optimizing over the set of separable quantum states

Section 2 showed that there always *exists* a product state assignment achieving a certain non-trivial approximation ratio. In this section, we show how to efficiently *find* such a product state. Our main theorem of this section is the following (Thm. 10), from which Thm. 4 follows easily (see discussion at end of Sec. 3.3).

Theorem 10. *Let H be an instance of MAX-k-local Hamiltonian acting on n qudits, and let* OPT_P *denote the optimum value of* $\mathrm{Tr}(H\rho)$ *over all* product *states* $\rho \in D((\mathbb{C}^d)^{\otimes n})$*. Then, for any fixed* $\epsilon > 0$*, there exists a polynomial time (deterministic) algorithm which outputs* $\rho_1 \otimes \cdots \otimes \rho_n \in D((\mathbb{C}^d)^{\otimes n})$ *such that* $\mathrm{Tr}(H\rho_1 \otimes \cdots \otimes \rho_n) \geq \mathrm{OPT}_P - \epsilon n^k$.

We first outline our approach by generalizing the discussion in Sec. 1, introducing tools and notation we will require along the way. The optimal value OPT_P over product state assignments for any MAX-k-local Hamiltonian instance can be expressed as the following program, denoted P_1:

$$\mathrm{OPT}_P \;=\; \max \sum_{i_1,\ldots,i_k}^{n} \mathrm{Tr}(H_{i_1,\ldots,i_k}\rho_{i_1} \otimes \cdots \otimes \rho_{i_k}) \;\text{ s.t. }\; \rho_i \succeq 0 \;\text{ and }\; \mathrm{Tr}(\rho_i) = 1 \;\text{ for } 1 \leq i \leq n. \tag{7}$$

As done in Eqn. (2), we now recursively decompose our objective function as a sequence of nested sums. Let $\{\sigma_i\}_{i=1}^d$ be a traceless, Hermitian orthogonal basis for the set of Hermitian operators acting on $\mathbb{C}^d$, such that $\mathrm{Tr}(\sigma_i\sigma_j) = 2\delta_{ij}$ (for δ_{ij} the Kroenecker delta) (see, e.g. [Kim03]). Then, by rewriting each $H_{i_1,\ldots,i_k}$ in terms of $\{\sigma_i\}_{i=1}^d$, our objective function becomes

$$\begin{aligned}\sum_{i_k,\ldots,i_1}^{n} \mathrm{Tr}\left[\left(\sum_{j_k,\ldots,j_1=1}^{d^2} r^{i_1,\ldots,i_k}_{j_1,\ldots,j_k}\sigma_{j_k} \otimes \cdots \otimes \sigma_{j_1}\right)\rho_{i_k} \otimes \cdots \otimes \rho_{i_1}\right] = \\ \sum_{i_k,j_k} \mathrm{Tr}(\sigma_{j_k}\rho_{i_k}) \left[\sum_{i_{k-1},j_{k-1}} \mathrm{Tr}(\sigma_{j_{k-1}}\rho_{i_{k-1}}) \left[\cdots \left[\sum_{i_1} \mathrm{Tr}\left(\left(\sum_{j_1} r^{i_1,\ldots,i_k}_{j_1,\ldots,j_k}\sigma_{j_1}\right)\rho_{i_1}\right)\right]\right]\right],\end{aligned} \tag{8}$$

where each $\mathbf{r}^{i_1,\ldots,i_k} \in \mathbb{R}^{d^2}$. We henceforth think of the objective function above as a "degree-k inner product", i.e. as a sequence of k nested sums involving inner products, in analogy to the degree-k polynomials of Ref. [AKK99]. In this sense, a degree-1 inner product would refer to only the innermost sums over i_1 and j_1, and a degree-k inner product would denote the entire expression in Eqn. (8). More formally, we denote[5] a degree-b inner product for $1 \leq b \leq k$ using map $t_b : H(\mathbb{C}^d)^{\times n} \mapsto \mathbb{R}$, defined such that[6] $t_b(\rho_1, \ldots, \rho_n) := \sum_{i_b,j_b} \mathrm{Tr}(\sigma_{j_b}\rho_{i_b}) \left[\cdots \left[\sum_{i_1} \mathrm{Tr}\left(\left(\sum_{j_1} r^{i_1,\ldots,i_k}_{j_1,\ldots,j_k}\sigma_{j_1}\right)\rho_{i_1}\right)\right]\right]$.

Our approach is to "linearize" the objective function of P_1 using exhaustive sampling and recursion to estimate its degree-$(k-1)$ inner products. To do so, we will require the Sampling Lemma.

Lemma 11 (Sampling Lemma [AKK99])**.** *Let (a_i) be a sequence of n real numbers with $|a_i| \leq M$ for all i, and let $f > 0$. If we choose a multiset of $s = g \log n$ of the a_i at random (with replacement), then their sum q satisfies* $\sum_i a_i - nM\sqrt{\frac{f}{g}} \leq q\frac{n}{s} \leq \sum_i a_i + nM\sqrt{\frac{f}{g}}$ *with probability at least* $1 - n^{-f}$.

The proof of Lemma 11 follows from a simple application of the Höffding bound [Höf64]. To use the Sampling Lemma in conjunction with exhaustive sampling, we will discretize the space of 1-qudit density operators using a δ-net $G \subseteq H(\mathbb{C}^d)$, such that for all $\rho \in D(\mathbb{C}^d)$, there exists $\sigma \in G$ such that $\| \rho - \sigma \|_{\mathrm{F}} \leq \delta$. We now show how to construct G.

---

[5] See the beginning of App. A for more elaborate notation used in the proofs of the claims of Sec. 3.

[6] Note that t_b implicitly depends on parameters $i_{b+1}, \ldots, i_k$ and $j_{b+1}, \ldots, j_k$.

To obtain G, we instead construct a δ-net for a subset of $H(\mathbb{C}^d)$ which *contains* $D(\mathbb{C}^d)$, namely the set[7] $\mathcal{A}(\mathbb{C}^d) := \left\{A \in H(\mathbb{C}^d) \mid \max_{i,j} |A(i,j)| \leq 1\right\}$. Creating a δ-net over $\mathcal{A}(\mathbb{C}^d)$ is simple: we cast a (δ/d)-net over the unit disk for each of the complex $d(d-1)/2$ matrix entries above the diagonal, and likewise over $[-1,1]$ for the entries on the diagonal. Letting m and n denote the minimum number of points required to create such (δ/d)-nets for each of the diagonal and off-diagonal entries, respectively, we have that $|G| = m^{\frac{d(d-1)}{2}} n^d$. For example, simple nets of size $m \approx d/\delta$ and $n \approx d^2/\delta^2$ can be obtained by placing a 1D and 2D grid over $[-1,1]$ and the length 2 square in the complex plane centered at $(0,0)$, respectively, implying $|G| \in O(1)$ when $d \in O(1)$. To show that G is indeed a δ-net, we now bound the Frobenius[8] distance between arbitrary $\rho \in D(\mathbb{C}^d)$ and the closest $\tilde{\rho} \in G$. Specifically, let $A := \rho - \tilde{\rho}$. Then:

$$\| A \|_{\mathrm{F}} = \sqrt{\mathrm{Tr}(A^\dagger A)} = \sqrt{\sum_{ij} |A(i,j)|^2} \leq \sqrt{\sum_{ij} (\delta/d)^2} = \frac{\delta}{d}(d) = \delta.$$

Finally, we remark that our *dense* assumption on MAX-k-local Hamiltonian instances is only necessary to convert the absolute error of Thm. 10 to a relative one [GK1] (this conversion is detailed in Sec. 3.3). The remaining sections are organized as follows: In Sec. 3.1, we show how to recursively estimate degree-b inner products using the Sampling Lemma. We then use this estimation technique in Sec. 3.2 to linearize our optimization problem P_1. Sec. 3.3 brings everything together by presenting and analyzing the complete approximation algorithm. To ease reading of the remaining sections, all technical proofs are found in App. A. Please see the beginning of App. A for definitions of the more elaborate notation used in these proofs.

3.1 Estimating degree-k inner products using the Sampling Lemma

Our recursive procedure, EVAL, for estimating a degree-k inner product using the Sampling Lemma is stated as Alg. 12. There are two sources of error we must analyze: the Sampling Lemma, and our δ-net over $\mathbb{C}^d$. We claim that EVAL estimates the degree-b inner product $t_b(\rho_1, \ldots, \rho_n)$ to within additive error $\pm\epsilon_b n^b$, where ϵ_b is defined as follows. Set $\Delta := \sqrt{2}d(1+\delta)$, for δ from our δ-net. Then,

$$\epsilon_b := d^{\frac{k}{2}} \left(\sqrt{\frac{f}{g}} + \delta \right) \left(\frac{\Delta^b - 1}{\Delta - 1} \right). \tag{9}$$

The following lemma formalizes this claim. We adopt the convention of [AKK99] and let $x \in y \pm z$ denote $x \in [y, z]$. Alg. 12 is our operator analogue of the algorithm *Eval* in Section 3.3 of [AKK99].

Lemma 13. *Let* $t_k : H(\mathbb{C}^k)^{\times n} \mapsto \mathbb{R}$ *be defined using set* $\left\{H_{i_1,\ldots,i_k}\right\} \subseteq H((\mathbb{C}^d)^{\otimes k})$ *(as in Eqn. (8)). Let* $S \subseteq \{1, \ldots, n\}$ *such that* $|S| = g \log n$ *have its elements chosen uniformly at random with replacement. Let* $\rho_1, \ldots, \rho_n \in D(\mathbb{C}^d)$ *be some assignment on all* n *qudits, and* $\{\tilde{\rho}_i : i \in S\}$ *a set of elements in our* δ*-net such that* $\| \rho_i - \tilde{\rho}_i \|_{\mathrm{F}} \leq \delta$ *for all* $i \in S$. *Then, for* $1 \leq b \leq k$, *with probability at least* $1 - d^{2b} n^{b-f}$, *we have* $\mathrm{EVAL}(t_b, S, \{\tilde{\rho}_i : i \in S\}) \in t_b(\rho_1, \ldots, \rho_n) \pm \epsilon_b n^b$, *where* ϵ_b *is defined as in Eqn. (9).*

3.2 Linearizing our optimization problem

Our procedure, LINEARIZE, for "linearizing" the objective function of P_1 using EVAL from Sec. 3.1 is stated as Alg. 14. Alg. 14 takes as input P_1 and a set of sample points $\{\tilde{\rho}_i\}$, and outputs a semidefinite program (SDP) which we shall henceforth refer to as P_2. We remark that LINEARIZE is our version of

---

[7] Note: A net over $\mathcal{A}(\mathbb{C}^d)$ may allow non-positive assignments for a qudit. See Sec. 3.3 for why this is of no consequence.

[8] We use the Frobenius norm as it allows a simple analysis. It is straightforward, however, to switch to say the trace norm using the fact that $\| X \|_{\mathrm{F}} \leq \sqrt{d} \| X \|_{\mathrm{tr}}$ for all $X \in \mathbb{C}^d$.

Algorithm 12. EVAL(t_b , S , $\{\tilde{\rho}_i : i \in S\}$).

```
• Input:  (1) A degree-$b$ inner product $t_b : H(\mathbb{C}^d)^{\times n} \mapsto \mathbb{R}$ for $1 \leq b \leq k$
          (2) A subset $S \subseteq \{1, \dots, n\}$ of size $|S| = O(\log n)$
          (3) Sample points $\{\tilde{\rho}_i : i \in S\}$ such that $\| \tilde{\rho}_i - \rho_i \|_{\mathrm{F}} \leq \delta$ for all $i \in S$
• Output: $x \in \mathbb{R}$ such that $x \in t_b(\rho_1, \dots, \rho_n) \pm \epsilon_b n^b$ (for $\epsilon_b$ defined in Eqn. (9)).
1. For all $i \in S$ and $j = 1 \dots d^2$:
   (a) (Base Case) if $b = 1$, set $e_{ij} = 1$.
   (b) (Recurse) else, set $e_{ij} = \mathrm{EVAL}(t_{b-1}^{ij}, S, \{\tilde{\rho}_i : i \in S\})$.
2. Return $\frac{n}{|S|} \sum_{i \in S} \left[ \sum_{j=1}^{d^2} \mathrm{Tr}(\sigma_j \tilde{\rho}_i) e_{ij} \right]$.
```

the procedure *Linearize* in Sec. 3.4 of [AKK99], extended to the setting of operators and a more complex error structure. Although LINEARIZE is presented as linearizing an objective function here, the same techniques straightforwardly apply in linearizing constraints involving high-degree inner products.

To prove correctness of our final approximation algorithm, we require the following two important lemmas regarding P_2. The first shows that any feasible solution $(\rho_1, \dots, \rho_n)$ for P_1 consistent with the sample set $\{\tilde{\rho}_i : i \in S\}$ fed into LINEARIZE is also a feasible solution for P_2 with high probability.

Lemma 15. *Let t_k, assignment $(\rho_1, \dots, \rho_n)$, S, and $\{\tilde{\rho}_i : i \in S\}$ be defined as in Lem. 13. Then, for any $f, g > 0$, calling LINEARIZE with parameters t_k, $\{\tilde{\rho}_i : i \in S\}$, and $\epsilon = \epsilon_k$ (for ϵ_k defined in Eqn. (9)) yields an SDP P_2 for which the assignment $\{\rho_1, \dots, \rho_n\}$ is feasible with probability at least $1 - d^{2k} n^{k-f}$.*

The second lemma is a bound on how far the optimal solution of P_2 is from the optimal solution for P_1. We adopt the convention of [AKK99] and write $[x, y] \pm z$ to denote interval $[x - z, y + z]$.

Lemma 16. *Let OPT_P be the optimal value for P_1 with corresponding assignment $\rho^{\mathrm{OPT}_P} := (\rho_1^{\mathrm{opt}}, \dots, \rho_n^{\mathrm{opt}})$, and let $\{\tilde{\rho}_i : i \in S\}$ be such that $\left\| \tilde{\rho}_i - \rho_i^{\mathrm{opt}} \right\|_{\mathrm{F}} \leq \delta$ for all $i \in S$ for some $S \subseteq \{1, \dots, n\}$. Let P_2 denote the SDP obtained by calling LINEARIZE with S, and denote by ϵ_m for $1 \leq m \leq k$ the error parameter passed with map t_m into a (possibly recursive) call to LINEARIZE. Then, letting OPT_2 denote the optimal value of P_2, we have with probability at least $1 - d^{2k} n^{k-f}$ (for parameters set as in Lem. 15) that $\mathrm{OPT}_2 \in \mathrm{OPT}_P \pm d(d + \sqrt{2}) \left[\sum_{m=1}^{k-1} (\sqrt{2} d)^{k-1-m} \epsilon_m \right] n^k$.*

3.3 The final algorithm

We finally present our approximation algorithm, APPROXIMATE (Alg. 17), in its entirety, which exploits our ability to linearize P_1 using LINEARIZE (Alg. 14). This proves Thm 10, which in turn implies Thm. 4. We first clarify a few points about APPROXIMATE, then analyze its runtime, and follow with further discussion, including the algorithm's derandomization and a proof that dense MAX-*k*-local Hamiltonian remains QMA-hard.

We begin by explaining the rationale behind the constants in Alg. 17. The constant $\varepsilon_{\mathrm{sdp}}$ is the additive error incurred when solving an SDP [GLS93]. We choose ϵ' so that after running LINEARIZE and solving P_2^i, the total additive error is at most ϵ, as desired. We choose f to ensure the probability of success is at least $1/2$. Finally, we set g large enough and δ (for our δ-net) small enough to ensure that ϵ' matches the error bounds for EVAL in Lem. 13.

Algorithm 14. LINEARIZE(t_b , $\mathcal{N}$, S, $\{\tilde{\rho}_i : i \in S\}$, ϵ, U, L).

```
• Input:  (1) A degree-$b$ inner product $t_b : H(\mathbb{C}^d)^{\times n} \mapsto \mathbb{R}$ for $1 \leq b \leq k$.
          (2) A set of linear constraints $\mathcal{N}$ (e.g. "$\rho_i \succeq 0$").
          (3) A subset $S \subseteq \{1, \ldots, n\}$ of size $|S| = O(\log n)$.
          (4) Sample points $\{\tilde{\rho}_i : i \in S\}$ consistent with some feasible solution $(\rho_1, \ldots, \rho_n)$ for
              $P_1$ such that $\| \tilde{\rho}_i - \rho_i \|_{\mathrm{F}} \leq \delta$ for all $i \in S$.
          (5) An error parameter $\epsilon > 0$.
          (6) (Optional) upper and lower bounds $U, L \in \mathbb{R}$. If $U$ and $L$ are not provided, we
              assume $U, L = \infty$.

• Output: (1) (Optional) A linear objective function $f : (L(\mathbb{C}^d))^{\times n} \to \mathbb{R}$.
          (2) An updated set of linear constraints, $\mathcal{N}$.

1. (Base case) If $b = 1$, then
   (a) (Trivial: Initial objective function was linear) If $U = L = \infty$, return $[t_b, \mathcal{N}]$.
   (b) (Reached bottom of recursion) Else, return $[\mathcal{N} \cup \{``L \leq t_b(\rho_1, \ldots, \rho_n) \leq U"\}]$.
2. (Recursive case) For $i = 1 \ldots n$ and $j = 1 \ldots d^2$ do
   (a) Set $e_{ij} := \mathrm{EVAL}(t_{b-1}^{ij}, S, \{\tilde{\rho}_i : i \in S\})$.
   (b) Set $\epsilon' := \epsilon - d^{\frac{k}{2}} \left( \sqrt{\frac{f}{g}} + \delta \right) \Delta^{b-1}$, for $\Delta$ defined in Eqn. (9).
   (c) Set $l_{ij} := e_{ij} - \epsilon' n^{b-1}$ and $u_{ij} := e_{ij} + \epsilon' n^{b-1}$.
   (d) Call LINEARIZE($t_{b-1}^{ij}, \mathcal{N}, S, \{\tilde{\rho}_i : i \in S\}, \epsilon', u_{ij}, l_{ij}$).
3. (a) (Entire computation done) If $U = L = \infty$, return $\left[\sum_{ij} \mathrm{Tr}(\sigma_j \rho_i) e_{ij}, \mathcal{N}\right]$.
   (b) (Recursive call done) Else, return $\left[\mathcal{N} \cup \left\{``L - \epsilon' d^2 n^b \leq \sum_{ij} \mathrm{Tr}(\sigma_j \rho_i) e_{ij} \leq U + \epsilon' d^2 n^b"\right\}\right]$.
```

We now analyze the runtime of Alg. 17. Let $|G|$ denote the size of our δ-net G for a qudit. Then, for each of the $|G|^{g \log n}$ iterations of line 6, we first take $O(n^{k-1})$ time to run LINEARIZE, outputting $O(n^{k-1})$ new linear constraints (seen via a simple inductive argument). We then solve SDP P_2^i, which can be done in time polynomial in n and $\log(1/\varepsilon_{\mathrm{sdp}})$ using the ellipsoid method [GLS93] (see, e.g., [Wat09]). Let $r(n, \varepsilon_{\mathrm{sdp}})$ denote the maximum runtime required to solve any of the P_2^i. Then, the overall runtime for Alg. 17 is $O(n^{g \log |G|}(n^{k-1} + r(n, \varepsilon_{\mathrm{sdp}})))$, which is polynomial in n for $\epsilon, d, k \in O(1)$ (recall from Sec. 3 that $|G| \in O((\frac{d}{\delta})^d)$, and that δ and g are constant in our setting). Note that, due to the implicit dependence of g on ϵ, this runtime scales at least exponentially with varying ϵ.

Before moving to further discussion, we make two remarks. First, one can convert the output of Alg. 17 to a *pure* state with the same guarantee by adapting the standard classical *method of conditional expectations* [Vaz01]. To demonstrate, suppose $\{\rho_i\}$ is output by Alg. 17. Then, set ρ_1' to be the eigenvector $|\psi_j\rangle\langle\psi_j|$ of ρ_1 for which the assignment $|\psi_j\rangle\langle\psi_j| \otimes \rho_2 \otimes \cdots \otimes \rho_n$ performs best[9] for P_1. Let our new assignment be $\rho_1' \otimes \rho_2 \otimes \cdots \otimes \rho_n$. Now repeat for each ρ_i for $2 \leq i \leq n$. The final state $\rho_1' \otimes \cdots \otimes \rho_n'$ is pure, and by convexity is guaranteed to perform as well as $\rho_1 \otimes \cdots \otimes \rho_n$.

Second, recall from Sec. 3 that we constructed a δ-net over a space larger than $D(\mathbb{C}^d)$, allowing possibly non-positive assignments for a qudit. We now see that this is of no consequence, since regardless of which samples (positive or not) we use to derive our estimates with the Sampling Lemma, any feasible solution to P_2^i in Alg. 17 is a valid assignment for P_1. Moreover, we know that for each optimal ρ_i for

---

[9] If the spectrum of ρ_i is degenerate, begin by fixing an arbitrary choice of spectral decomposition for ρ_i.

Algorithm 17. APPROXIMATE(H, ϵ).

```
• Input: (1) A k-local Hamiltonian $H = \sum_{i_1,\dots,i_k} H_{i_1,\dots,i_k}$ for each $H_{i_1,\dots,i_k} \in H((\mathbb{C}^d)^k)$.
         (2) An error parameter $\epsilon > 0$.
• Output: A product assignment $\rho_1 \otimes \cdots \otimes \rho_n$ that with probability at least $1/2$, has value at least
          $\mathrm{OPT}_P - \epsilon n^k$, for $\mathrm{OPT}_P$ the optimal value for $H$ over all product state assignments.
1. Set $\varepsilon_{\mathrm{sdp}} := \varepsilon/10$.
2. Define $h : \mathbb{R} \rightarrow \mathbb{R}$ such that for any error parameter $\epsilon$ input to LINEARIZE, $h(\varepsilon)n^k$ is the
   absolute value of the bound on additive error given by Lem. 16. Then, define $\epsilon'$ implicitly so that
   $h(\epsilon') + \varepsilon_{\mathrm{sdp}} = \epsilon$ holds.
3. Define constant $f$ such that $1 - d^{2k}n^{k-f} > 1/2$.
4. Define constants $g$ and $\delta$ implicitly so that $\epsilon' = d^{\frac{k}{2}}\left(\sqrt{\frac{f}{g}} + \delta\right)\left(\frac{\Delta^k-1}{\Delta-1}\right)$, for $\Delta$ defined in Eqn. (9).
5. Choose $g \log n$ indices $S \subseteq \{1,\dots,n\}$ independently and uniformly at random.
6. For each possible assignment $i$ from our $\delta$-net to the qudits in $S$:
   (a) Call LINEARIZE($t_k$, {$P_1$'s constraints}, $S$, $i$, $\epsilon'$) to obtain SDP $P_2^i$.
   (b) Let $\alpha_i$ denote the value of $P_1$ obtained by substituting in the optimal solution of $P_2^i$.
7. Return the assignment corresponding to the maximum over all $\alpha_i$.
```

P_1, there must be *some* operator (positive or not) within distance δ in our net, ensuring our estimates obtained using the Sampling Lemma are within our error bounds.

Converting the absolute error of Algorithm 17 into relative error. To convert the absolute error $\pm\epsilon n^k$ of Alg. 17 into a *relative* error of $1-\epsilon'$ for any ϵ', define constant c such that cn^k is the value obtained for a MAX-k-local Hamiltonian instance by choosing the maximally mixed assignment I/d^n (analogous to a classical random assignment). Since I/d^n can be written as a mixture of computational basis states, we have $\mathrm{OPT}_P \geq cn^k$. It follows that by setting $\epsilon = c\epsilon'$, Alg. 17 returns an assignment with value at least $\mathrm{OPT}_P - c\epsilon' n^k \geq \mathrm{OPT}_P - \epsilon'\mathrm{OPT}_P \geq \mathrm{OPT}_P(1-\epsilon')$, as desired.

Derandomizing Algorithm 17. The source of randomness in our algorithm is Lem. 11. By a standard argument in [AKK99] (see also [BR94, BGG93]), this randomness can be eliminated with only polynomial overhead. Specifically, we replace the random selection of $g \log n$ indices in the Sampling Lemma with the set of indices encountered on a random walk of length $O(g \log n)$ along a constant degree expander [Gil93]. Since the expander has constant degree, we can efficiently deterministically iterate through all $n^{O(g)}$ such walks, and since such a walk works with probability $1/n^{O(1)}$, at least one walk will work for all $\mathrm{poly}(n)$ sampling experiments we wish to run.

QMA-hardness of dense MAX-*k*-local Hamiltonian. It is easy to see that (exact) MAX-2-local Hamiltonian remains QMA-hard for dense instances (a similar statement holds for MAX-2-SAT [AKK99]). For any MAX-2-local Hamiltonian instance with optimal value OPT, we simply add n qudits, between any two of which we place the constraint $|00\rangle\langle 00|$ (no constraints are necessary between old and new qudits). Then, the new Hamiltonian has optimal value $\mathrm{OPT} + \binom{n}{2}$, making it dense, and the ability to solve this new instance implies the ability to solve the original one. The argument extends straightforwardly to MAX-k-local Hamiltonian for $k > 2$.

4 Acknowledgements

We thank Jamie Sikora and Sarvagya Upadhyay for helpful feedback on an earlier version of this draft, and Yi-Kai Liu for interesting discussions. We wish to especially thank Oded Regev for many helpful comments and suggestions, and Richard Cleve for bringing our attention to the method of conditional expectations, and for stimulating discussions and support.

References

[AAI10] D. Aharonov, I. Arad, and S. Irani. Efficient algorithm for approximating one-dimensional ground states. *Physical Review A*, 82:012315, 2010.

[AALV09] D. Aharonov, I. Arad, Z. Landau, and U. Vazirani. The detectibility lemma and quantum gap amplification. In *Proceedings of the 41st annual ACM syposium on Theory of computing*, volume 287, pages 417–426, 2009.

[Aar06] S. Aaronson. The quantum PCP manifesto. http://scottaaronson.com/blog/?p=139, 2006.

[AB09] S. Arora and B. Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, 2009.

[AdlVKK02] N. Alon, W. F. de la Vega, R. Kannan, and M. Karpinski. Random sampling and approximation of MAX-CSP problems. In *Proceedings of the 34th Symposium on Theory of computing*, pages 232–239, 2002.

[AGIK09] D. Aharonov, D. Gottesman, S. Irani, and J. Kempe. The power of quantum systems on a line. *Communications in Mathematical Physics*, 287:41–65, 2009.

[AKK99] S. Arora, D. Karger, and M. Karpinski. Polynomial time approximation schemes for dense instances of NP-hard problems. *Journal of Computer and System Sciences*, 58:193–210, 1999.

[ALM$^+$98] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, 1998.

[AM08] P. Austrin and E. Mossel. Approximation resistant predicates from pairwise independence. In *Proceedings of the 23rd IEEE Conference on Computational Complexity*, pages 249–258, 2008.

[AS98] S. Arora and S. Safra. Probabilistic checking of proofs: A new characterization of NP. *Journal of the ACM*, 45(1):70–122, 1998.

[AvDK$^+$07] D. Aharanov, W. van Dam, J. Kempe, Z. Landau, S. Lloyd, and O. Regev. Adiabatic quantum computation is equivalent to standard quantum computation. *SIAM Journal of Computing*, 37(1):166–194, 2007.

[BBT09] N. Bansal, S. Bravyi, and B. M. Terhal. Classical approximation schemes for the ground-state energy of quantum and classical Ising spin Hamiltonians on planar graphs. *Quantum Information and Computation*, 9(7&8):0701–0720, 2009.

[BdlVK03] C. Bazgan, W. F. de la Vega, and M. Karpinski. Polynomial time approximation schemes for dense instances of minimum constraint satisfaction. *Random Structures & Algorithms*, 23(1):73–91, 2003.

[BDOT08] S. Bravyi, D. P. Divincenzo, R. Oliveira, and B. M. Terhal. The complexity of the stoquastic local Hamiltonian problems. *Quantum Information and Computation*, 8(5):0361–0385, 2008.

[BGG93] M. Bellare, O. Goldreich, and S. Goldwasser. Randomness in interactive proofs. *Computational Complexity*, 3:319—354, 1993.

[BR94] M. Bellare and J. Rompel. Randomness-efficient oblivious sampling. In *Proceedings of the 28th Symposium on the Foundations of Computer Science*, pages 276—287, 1994.

[Bra06] S. Bravyi. Efficient algorithm for a quantum analogue of 2-SAT. *arXiv:quant-ph/0602108v1*, 2006.

[BS07] S. Beigi and P. W. Shor. On the complexity of computing zero-error and Holevo capacity of quantum channels. *arXiv:0709.2090v3*, 2007.

[BV05] S. Bravyi and M. Vyalyi. Commutative version of the local Hamiltonian problem and common eigenspace problem. *Quantum Information and Computation*, 5(3):187–215, 2005.

[CMM07] M. Charikar, K. Makarychev, and Y. Makarychev. Near-optimal algorithms for maximum constraint satisfaction problems. In *Lecture Notes in Computer Science*, volume 4627, pages 149–163, 2007.

[CV09] J. I. Cirac and F. Verstraete. Renormalization and tensor product states in spin chains and lattices. *Journal of Physics A: Mathematical and Theoretical*, 42(50):504004, 2009.

[dlV96] W. F. de la Vega. MAX-CUT has a randomized approximation scheme in dense graphs. *Random Structures & Algorithms*, 8(3):187–198, 1996.

[dlVK00] W. F. de la Vega and M. Karpinski. Polynomial time approximation of dense weighted instances of MAX-CUT. *Random Structures & Algorithms*, 16:314–332, 2000.

[dlVKKV05] W. F. de la Vega, M. Karpinski, R. Kannan, and S. Vempala. Tensor decomposition and approximation schemes for constraint satisfaction problems. In *Proceedings of the 37th Symposium on Theory of computing*, pages 747–754. ACM Press, 2005.

[DPS04] A. C. Doherty, P. A. Parrilo, and F. M. Spedalieri. Complete family of separability criteria. *Physical Review A*, 69:022308, 2004.

[FK96] A. M. Frieze and R. Kannan. The regularity lemma and approximation schemes dense problems. In *Proceedings of the 37th Annual Symposium on Foundations of Computer Science*, pages 12–20, 1996.

[GGR98] O. Goldreich, S. Goldwasser, and D. Ron. Property testing and its connection to learning and approximation. *Journal of the ACM*, 45(4):653–750, 1998.

[Gil93] D. Gillman. A Chernoff bound for random walks on expanders. In *Proceedings of the 34th Symposium on the Foundations of Computer Science*, pages 680—691, 1993.

[GK1] We briefly discuss why it is not necessary to have a *dense* MAX-k-local Hamiltonian instance in order to apply the Sampling Lemma: Specifically, observe that the Sampling Lemma assumes there are n terms in the sum to be estimated, and that we are able to determine s of them. Looking back at Eqn. (1) and considering, say, qudit i, if we wish to use the Sampling Lemma to estimate the inner sum over neighbours $N(i)$ of i, we might run into a problem if i does *not* have $\Theta(n)$ neighbours. To circumvent this [AKK99], observe that Lem. 11 only gives us an estimate to within $\pm\epsilon n$. Thus, if $N(i) \leq \epsilon n/10$ (say), then we do not use the Sampling Lemma, but rather let our estimate be simply 0, which is guaranteed to fall within the desired error bounds (observe an estimate of 0 does not necessarily work, on the other hand, if $N(i)$ is large (say $N(i) = n - 1$), since typically $f/g < 1$). Throughout the remainder of our discussion, we assume this cutoff principle is implicitly present when employing Lem. 11.

[GLS93] M. Grötschel, L. Lovàsz, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Springer-Verlag, 1993.

[Hås97] J. Håstad. Some optimal inapproximability results. In *Proceedings of the 29th Symposium on Theory of computing*, pages 1–10, 1997.

[Has05] G. Hast. Approximating Max kCSP - outperforming a random assignment with almost a linear factor. In *Proceedings of the 32nd International Colloquium on Automata, Languages, and Programming*, pages 956–968, 2005.

[Hås07] J. Håstad. On the approximation resistance of a random predicate. In *Lecture Notes in Computer Science*, volume 4627, pages 149–163, 2007.

[Hoc97] D. Hochbaum. *Approximation Algorithms for NP-Hard Problems*. Wadsworth Publishing Company, 1997.

[Höf64] W. Höffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1964.

[JGL10] S. P. Jordan, D. Gosset, and P. J. Love. QMA-complete problems for stoquastic Hamiltonians and Markov matrices. *Physical Review A*, 81(3):032331, 2010.

[JJUW10] R. Jain, Z. Ji, S. Upadhyay, and J. Watrous. QIP=PSPACE. In *Proceedings of the 42nd Symposium on Theory of Computing*, pages 573—581, 2010.

[Kho02] S. Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the 34th Symposium on Theory of computing*, pages 767—775, 2002.

[Kim03] G. Kimura. The Bloch vector for N-level systems. *Physics Letters A*, 314(5), 2003.

[KKR06] J. Kempe, A. Kitaev, and O. Regev. The complexity of the local Hamiltonian problem. *SIAM Journal of Computing*, 35(5):1070–1097, 2006.

[KR03] J. Kempe and O. Regev. 3-local Hamiltonian is QMA-complete. *Quantum Information and Computation*, 3(3):258–264, 2003.

[KSTW01] S. Khanna, M. Sudan, L. Trevisan, and D. Williamson. The approximability of constraint satisfaction problems. *SIAM Journal of Computing*, 30(6):1863–1920, 2001.

[KSV02] A. Kitaev, A. Shen, and M. Vyalyi. *Classical and Quantum Computation*. American Mathematical Society, 2002.

[LCV07] Y.-K. Liu, M. Christandl, and F. Verstraete. Quantum computational complexity of the N-representability problem: QMA complete. *Physical Review Letters*, 98:110503, 2007.

[Liu06] Y.-K. Liu. Consistency of local density matrices is QMA-complete. In *Lecture Notes in Computer Science*, volume 4110, pages 438–449, 2006.

[LLM$^+$10] C. .R. Laumann, A. M. Läuchli, R. Moessner, A. Scardicchio, and S. L. Sondhi. Product, generic, and random generic quantum satisfiability. *Physical Review A*, 81:062345, 2010.

[LLZ02] M. Lewin, D. Livnat, and U. Zwick. Improved rounding techniques for MAX 2-SAT and MAX DI-CUT problems. In *Proceedings of the 9th International IPCO Conference on Integer Programming and Combinatorial Optimization*, pages 67–82, 2002.

[LMRS10] T. Lee, R. Mittal, B. W. Reichardt, and R. Spalek. An adversary for algorithms. *arXiv:1011.3020v1*, 2010.

[ÖR95] S. Östlund and S. Rommer. Thermodynamic limit of density matrix renormalization. *Physical Review Letters*, 75:3537–3540, 1995.

[OT08] R. Oliveira and B. M. Terhal. The complexity of quantum spin systems on a two-dimensional square lattice. *Quantum Information and Computing*, 8(10):0900–0924, 2008.

[PWKE98] I. Peschel, X. Wang, M. Kaulke, and K. Hallberg (Edgs.). Density-matrix renormalization - a new numerical method in physics. In *Lecture Notes in Physics*, volume 528. Springer-Verlag, 1998.

[RÖ97] S. Rommer and S. Östlund. Class of ansatz wave functions for one-dimensional spin systems and their relation to the density matrix renormalization group. *Physical Review B*, 55:2164–2181, 1997.

[Ros09] B. Rosgen. Testing non-isometry is QMA-complete. *arXiv:0910.3740v2*, 2009.

[SC10] N. Schuch and J. I. Cirac. Matrix product state and mean-field solutions for one-dimensional systems can be found efficiently. *Physical Review A*, 82:012314, 2010.

[Sch05] U. Schollwöck. The density-matrix renormalization group. *Reviews of Modern Physics*, 77:259–315, 2005.

[ST06] A. Samorodnitsky and L. Trevisan. Gowers uniformity, influences of variables, and pcps. In *Proceedings of the 38th Symposium on Theory of computing*, pages 11–20, 2006.

[SV09] N. Schuch and F. Verstraete. Computational complexity of interacting electrons and fundamental limitations of density functional theory. *Nature Physics*, 5:732–735, 2009.

[Tre98] L. Trevisan. Parallel approximation algorithms by positive linear programming. *Algorithmica*, 21(1):72–88, 1998.

[TWP09] S. Tamaryan, T.-C. Wei, and D. Park. Maximally entangled three-qubit states via geometric measure of entanglement. *Physical Review A*, 80:052315, 2009.

[Vaz01] V. Vazirani. *Approximation Algorithms*. Springer, 2001.

[Wat09] J. Watrous. Semidefinite programs for completely bounded norms. *Theory of Computing*, 5:217—238, 2009.

[WG03] T.-C. Wei and P. M. Goldbart. Geometric measure of entanglement and applications to bipartite and multipartite quantum states. *Physical Review A*, 68:042307, 2003.

[Whi92] S. R. White. Density matrix formulation for quantum renormalization groups. *Physical Review Letters*, 69:2863–2866, 1992.

[Whi93] S. R. White. Density-matrix algorithms for quantum renormalization groups. *Physical Review B*, 48:10345–10356, 1993.

[WMN10] T.-C. Wei, M. Mosca, and A. Nayak. Interacting boson problems can be QMA hard. *Physical Review Letters*, 104:040501, 2010.

A Technical proofs for Section 3

Expanded Notation. We now expand on our previous notation for analyzing Eqn. (8) in order to facilitate proofs of the claims in Sec. 3. First, to recursively analyze a clause $H_{i_1,\ldots,i_k} \subseteq H((\mathbb{C}^d)^{\otimes k})$, let $H_b \in H((\mathbb{C}^d)^{\otimes b})$ for any $1 \leq b \leq k$ denote the action of $H_{i_1,\ldots,i_k}$ restricted to the first b of its k target qubits, i.e. $H_b := \sum_{j_b,\ldots,j_1}^{d^2} r^{i_1,\ldots,i_k}_{j_1,\ldots,j_k}\sigma_{j_b} \otimes \cdots \otimes \sigma_{j_1}$. For example, $H_1 = \sum_{j_1}^{d^2} r^{i_1,\ldots,i_k}_{j_1,\ldots,j_k}\sigma_{j_1}$ and $H_k = H_{i_1,\ldots,i_k}$. Note that H_b implicitly depends on variables $i_1,\ldots,i_k, j_{b+1},\ldots,j_k$, but to reduce clutter, our notation does not explicitly denote this dependence unless necessary. Next, to recursively analyze a degree-a inner product, we define $t_{a,b} : H(\mathbb{C}^d)^{\times n} \mapsto \mathbb{R}$ for any $0 \leq a \leq k$ and $1 \leq b \leq k$ such that $t_{a,b}(\rho_1,\ldots,\rho_n) := \sum_{i_a,\ldots,i_1}^{n} \operatorname{Tr}\left(H_b^{i_1,\ldots,i_k}\rho_{i_b} \otimes \cdots \otimes \rho_{i_1}\right)$ (where setting $a = 0$ eliminates the sum over indices i). For example, $t_{k,k}$ is our full "degree-k" objective function in Eqn. (7), and more generally, $t_{b,b}$ is the degree-b inner product in Eqn. (8). Allowing different values for a and b greatly eases our technical analysis. We use the shorthand t_b to denote $t_{b,b}$, and again only explicitly denote the dependence of $t_{a,b}$ on parameters $i_{a+1},\ldots,i_k$ and $j_{b+1},\ldots,j_k$ when necessary.

Lemma 18. *Let* $\{\rho_i\}_{i=1}^n \subseteq H(\mathbb{C}^d)$. *For any MAX-k-local Hamiltonian instance* $\left\{H_{i_1,\ldots,i_k}\right\} \subseteq H(\mathbb{C}^{d^k})$ *with decomposition for the* $H_{i_1,\ldots,i_k}$ *as given in Eqn. (8), we have for any* $0 \leq a \leq k$ *and* $1 \leq b \leq k$ *that* $|t_{a,b}(\rho_1,\ldots,\rho_n)| \leq \left(\max_{i_b,\ldots,i_1} \|\rho_{i_b}\|_{\mathrm{F}} \cdots \|\rho_{i_1}\|_{\mathrm{F}}\right) d^{\frac{k}{2}} n^a$.

Proof of Lem. 18. By the triangle inequality and the Hölder inequality for Schatten p-norms, we have

$$\begin{aligned} |t_{a,b}| = \left|\sum_{i_a,\ldots,i_1}^{n} \operatorname{Tr}\left(H_b\rho_{i_b} \otimes \cdots \otimes \rho_{i_1}\right)\right| \quad &\leq \quad \sum_{i_a,\ldots,i_1}^{n} \|H_b\|_{\mathrm{F}} \left\|\rho_{i_b} \otimes \cdots \otimes \rho_{i_1}\right\|_{\mathrm{F}} \\ &\leq \quad \left(\max_{i_b,\ldots,i_1} \|\rho_{i_b}\|_{\mathrm{F}} \cdots \|\rho_{i_1}\|_{\mathrm{F}}\right) \sum_{i_a,\ldots,i_1}^{n} \|H_b\|_{\mathrm{F}}, \end{aligned}$$

where we have used the fact that $\|A \otimes B\|_{\mathrm{F}} = \|A\|_{\mathrm{F}} \|B\|_{\mathrm{F}}$ for all $A, B \in L(\mathbb{C}^d)$. If we can now show that $\|H_b\|_{\mathrm{F}} \leq \|H_k\|_{\mathrm{F}}$ for all $1 \leq b \leq k$, then we would be done since we would have $\sum_{i_a,\ldots,i_1}^{n} \|H_b\|_{\mathrm{F}} \leq \|H_k\|_{\mathrm{F}} n^a \leq d^{\frac{k}{2}} n^a$, where $\|H_k\|_{\mathrm{F}} \leq d^{\frac{k}{2}}$ since $\|H_k\|_\infty \leq 1$ by definition. Indeed, we claim that for any fixed $1 \leq b \leq k$, we have $\|H_b\|_{\mathrm{F}} \leq 2^{\frac{b}{2}} \|H_k\|_{\mathrm{F}}$. To see this, note by straightforward expansion of the Frobenius norm and the fact that $\operatorname{Tr}(\sigma_i\sigma_j) = 2\delta_{ij}$ that

$$\|H_b\|_{\mathrm{F}} = \sqrt{\operatorname{Tr}(H_b^2)} = 2^{\frac{b}{2}} \sqrt{\sum_{j_b,\ldots,j_k} (r^{i_1,\ldots,i_k}_{j_1,\ldots,j_k})^2} \leq 2^{\frac{b}{2}} \left\|\mathbf{r}^{i_1,\ldots,i_k}\right\|_2 = 2^{\frac{b-k}{2}} \left(2^{\frac{k}{2}} \left\|\mathbf{r}^{i_1,\ldots,i_k}\right\|_2\right),$$

where $\mathbf{r}^{i_1,\ldots,i_k}$ is the coordinate vector of $H_{i_1,\ldots,i_k}$ from Eqn. (8). By the second equality in the chain above, we see that in fact $\|H_k\|_{\mathrm{F}} = 2^{\frac{k}{2}} \left\|\mathbf{r}^{i_1,\ldots,i_k}\right\|_2$, completing the proof of our claim. $\square$

***Proof of Lem. 13*.** We first derive the error bound of ϵ_b, and subsequently prove the probability bound. We follow [AKK99], and proceed by induction on b. For the base case $b = 1$, $\mathrm{EVAL}(H_1, S, \{\tilde{\rho}_i : i \in S\})$ attempts to estimate

$$t_1(\rho_1, \ldots, \rho_n) = \sum_{i_1} \left[\sum_{j_1} r^{i_1,\ldots,i_k}_{j_1,\ldots,j_k} \mathrm{Tr}(\sigma_{j_1} \rho_{i_1}) \right]$$

using our flawed sample points $\{\tilde{\rho}_i : i \in S\}$. To analyze the error of its output, assume first that our sample points are exact, i.e. $\tilde{\rho}_i = \rho_i$ for all $i \in S$. Then, by setting "a_i" in Lem. 11 to $t^{i_1}_{0,1}$ for $i = i_1$, and by using Lem. 18 with parameters $a = 0$ and $b = 1$ to obtain upper bound $M = d^{\frac{k}{2}}$, we have by the Sampling Lemma that (with probability at least $1 - n^{-f}$)

$$\frac{n}{|S|} \sum_{i_1 \in S} \left[\sum_{j_1} r^{i_1,\ldots,i_k}_{j_1,\ldots,j_k} \mathrm{Tr}(\sigma_{j_1} \rho_{i_1}) \right] \in t_1(\rho_1, \ldots, \rho_n) \pm d^{\frac{k}{2}} \sqrt{\frac{f}{g}}\, n. \tag{10}$$

This bound holds if we sum over exact sample points. If we instead sum over flawed sample points $\{\tilde{\rho}_i : i \in S\}$, the additional error is bounded by $\frac{n}{|S|}$ times

$$\left| \sum_{i_1 \in S} \left[\sum_{j_1} r^{i_1,\ldots,i_k}_{j_1,\ldots,j_k} \mathrm{Tr}(\sigma_{j_1} (\rho_{i_1} - \tilde{\rho}_{i_1})) \right] \right| \leq \sum_{i_1 \in S} \left| \sum_{j_1} r^{i_1,\ldots,i_k}_{j_1,\ldots,j_k} \mathrm{Tr}(\sigma_{j_1} (\rho_{i_1} - \tilde{\rho}_{i_1})) \right| \leq \sum_{i_1 \in S} (\| \rho_{i_1} - \tilde{\rho}_{i_1} \|_{\mathrm{F}}\, d^{\frac{k}{2}}) \leq d^{\frac{k}{2}} \delta n, \tag{11}$$

where the second inequality uses Lem. 18 with parameters $a = 0$ and $b = 1$ and the promise of our δ-net. We conclude for the base case that

$$\mathrm{EVAL}(H_1, S, \{\tilde{\rho}_i : i \in S\}) = \frac{n}{|S|} \sum_{i_1 \in S} \left[\sum_{j_1} r^{i_1,\ldots,i_k}_{j_1,\ldots,j_k} \mathrm{Tr}(\sigma_{j_1} \tilde{\rho}_{i_1}) \right] \in t_1(\rho_1, \ldots, \rho_n) \pm d^{\frac{k}{2}} \left(\sqrt{\frac{f}{g}} + \delta \right) n,$$

as desired.

Assume now that the inductive hypothesis holds for $1 \leq m \leq b - 1$. We prove the claim for $m = b$. To do so, suppose first that the recursive calls on line 1(b) of Alg. 12 return the *exact* values of $t^{ij}_{b-1}(\rho_1, \ldots, \rho_n)$, and that we have exact samples $\{\rho_i : i \in S\}$. Then, since by calling Lem. 18 with $a = b - 1$ we have $\left| \sum_j \mathrm{Tr}(\sigma_j \rho_i) t^{ij}_{b-1}(\rho_1, \ldots, \rho_n) \right| \leq d^{\frac{k}{2}} n^{b-1}$, it follows by the Sampling Lemma that

$$\frac{n}{|S|} \sum_{i \in S} \left[\sum_j \mathrm{Tr}(\sigma_j \rho_i) t^{ij}_{b-1}(\rho_1, \ldots, \rho_n) \right] \in \sum_{i=1}^{n} \left[\sum_j \mathrm{Tr}(\sigma_j \rho_i) t^{ij}_{b-1}(\rho_1, \ldots, \rho_n) \right] \pm d^{\frac{k}{2}} \sqrt{\frac{f}{g}}\, n^b. \tag{12}$$

To first adjust for using flawed samples, observe that an analogous calculation to Eqn. (11) yields $\left| \frac{n}{|S|} \sum_{i \in S} \left[\sum_j \mathrm{Tr}(\sigma_j (\rho_i - \tilde{\rho}_i)) \right] \right| \leq d^{\frac{k}{2}} \delta n^b$, where we have called Lem. 18 with $a = b - 1$. Thus, using flawed samples, the output of Alg. 12 satisfies

$$\frac{n}{|S|} \sum_{i \in S} \left[\sum_j \mathrm{Tr}(\sigma_j \tilde{\rho}_i) t^{ij}_{b-1} \right] \in \sum_{i=1}^{n} \left[\sum_j \mathrm{Tr}(\sigma_j \rho_i) t^{ij}_{b-1} \right] \pm d^{\frac{k}{2}} \left(\sqrt{\frac{f}{g}} + \delta \right) n^b. \tag{13}$$

To next drop the assumption that our estimates e_{ij} on line 1(b) are exact, apply the induction hypothesis to conclude that $e_{ij} \in t^{ij}_{b-1}(\rho_1, \ldots, \rho_n) \pm \epsilon_{b-1} n^{b-1}$. Then,

$$\begin{aligned} \frac{n}{|S|} \sum_{i \in S} \left[\sum_j \mathrm{Tr}(\sigma_j \tilde{\rho}_i) e_{ij} \right] &\in \frac{n}{|S|} \sum_{i \in S} \left[\sum_j \mathrm{Tr}(\sigma_j \tilde{\rho}_{ij}) \left(t^{ij}_{b-1} \pm \epsilon_{b-1} n^{b-1} \right) \right] \\ &\subseteq \frac{n}{|S|} \sum_{i \in S} \left[\sum_j \mathrm{Tr}(\sigma_j \tilde{\rho}_i) t^{ij}_{b-1} \right] \pm \frac{\epsilon_{b-1} n^b}{|S|} \sum_{i \in S} \left[\sum_{j=1}^{d^2} \mathrm{Tr}(\sigma_j \tilde{\rho}_i) \right] \\ &\subseteq \frac{n}{|S|} \sum_{i \in S} \left[\sum_j \mathrm{Tr}(\sigma_j \tilde{\rho}_i) t^{ij}_{b-1} \right] \pm \epsilon_{b-1} \sqrt{2} d (1 + \delta) n^b, \end{aligned} \tag{14}$$

where the last statement follows since

$$\left| \sum_{j=1}^{d^2} \mathrm{Tr}(\sigma_j \tilde{\rho}_i) \right| = \left| \sum_{j=1}^{d^2} \mathrm{Tr} \left(\sigma_j \left(\sum_{m=1}^{d^2} \tilde{r}_m \sigma_m \right) \right) \right| \leq 2 \sum_{m=1}^{d^2} |\tilde{r}_m| \leq 2d \parallel \tilde{\mathbf{r}} \parallel_2 \leq \sqrt{2} d (1 + \delta), \tag{15}$$

where $\tilde{\mathbf{r}}$ denotes the coordinate vector of $\tilde{\rho}_i$ with respect to basis $\{\sigma_m\}$, and we have used the facts that $\mathrm{Tr}(\sigma_i \sigma_j) = 2\delta_{ij}$, that $\parallel \mathbf{x} \parallel_1 \leq \sqrt{d} \parallel \mathbf{x} \parallel_2$ for $\mathbf{x} \in \mathbb{C}^d$, that $\parallel \tilde{\rho}_i \parallel_{\mathrm{F}} = \sqrt{2} \parallel \tilde{\mathbf{r}} \parallel_2$ for any $\tilde{\rho}_i \in H(\mathbb{C}^d)$, and that $\parallel \tilde{\rho}_i \parallel_{\mathrm{F}} \leq 1 + \delta$ (which follows from our δ-net and the triangle inequality). Thus, recalling that $\Delta = \sqrt{2} d (1 + \delta)$ and substituting Eqn. (13) into Eqn. (14), we have that

$$\frac{n}{|S|} \sum_{i \in S} \left[\sum_j \mathrm{Tr}(\sigma_j \tilde{\rho}_i) e_{ij} \right] \in t_b(\rho_1, \ldots, \rho_n) \pm \left[d^{\frac{k}{2}} \left(\sqrt{\frac{f}{g}} + \delta \right) + \epsilon_{b-1} \Delta \right] n^b.$$

We hence have the recurrence relation $\epsilon_b \leq d^{\frac{k}{2}} \left(\sqrt{\frac{f}{g}} + \delta \right) + \epsilon_{b-1} \Delta$, which when unrolled yields

$$\epsilon_b \leq d^{\frac{k}{2}} \left(\sqrt{\frac{f}{g}} + \delta \right) \sum_{m=0}^{b-1} \Delta^m = d^{\frac{k}{2}} \left(\sqrt{\frac{f}{g}} + \delta \right) \left(\frac{\Delta^b - 1}{\Delta - 1} \right),$$

as desired. This concludes the proof of the error bound.

To prove the probability bound, we instead prove the stronger bound of $1 - \left(\sum_{m=0}^{b-1} d^{2m} n^m \right) n^{-f}$ by induction on b. The base case $b = 1$ follows directly from our application of the Sampling Lemma in Eqn. (10). For the inductive step, define for brevity of notation $\gamma := d^2 n$, and apply the induction hypothesis to line 1(b) of Alg. 12 to conclude that each of the γ calls to EVAL fails will probability at most $(\sum_{m=0}^{b-2} \gamma^m) n^{-f}$. Then, by the union bound, the probability that at least one call fails is at most $(\sum_{m=1}^{b-1} \gamma^m) n^{-f}$. Similarly, since our application of the Sampling Lemma in line 2 of Alg. 12 fails with probability at most n^{-f}, we arrive at our claimed stronger bound of $1 - \left(\sum_{m=0}^{b-1} \gamma^m \right) n^{-f}$, as desired. ☐

Proof of Lem. 15. We begin by observing that if one sets $\epsilon = \epsilon_k$, then the value of ϵ' in line 2(b) of Alg. 14 is precisely ϵ_{k-1}, and more generally, the ϵ passed into the recursive call of line 2(e) on t_b for any $1 \leq b \leq k$ is ϵ_b. Now, focus on some recursive call on t_b for $b > 1$ (the case of $b = 1$ is straightforward by Lem. 13). If the estimates e_{ij} in line 2(a) succeed, then by Lem. 13, we know that $e_{ij} \in t^{ij}_{b-1}(\rho_1, \ldots, \rho_n) \pm \epsilon_{b-1} n^{b-1}$, implying $t^{ij}_{b-1}(\rho_1, \ldots, \rho_n) \in [l_{ij}, u_{ij}]$. Now, l_{ij} and u_{ij} are only incorporated into linear constraints in recursive calls on t^{ij}_{b-1}, yielding constraints of the form

$$l_{i_b j_b} - \epsilon_{b-2} d^2 n^{b-1} \leq \sum_{i_{b-1}, j_{b-1}} \mathrm{Tr}(\sigma_{j_{b-1}} \rho_{i_{b-1}}) e_{i_{b-1} j_{b-1}} \leq u_{i_b j_b} + \epsilon_{b-2} d^2 n^{b-1}. \tag{16}$$

But $\{\rho_1, \ldots, \rho_n\}$ must now satisfy this constraint, since recall

$$t_{b-1}(\rho_1, \ldots, \rho_n) = \sum_{i_{b-1}, j_{b-1}} \mathrm{Tr}(\sigma_{j_{b-1}} \rho_{i_{b-1}}) t_{b-2}^{i_{b-1} j_{b-1}}(\rho_1, \ldots, \rho_n),$$

and there are $d^2 n$ terms $e_{i_{b-1} j_{b-1}}$ in Eqn. (16) each yielding an additional error of at most $\epsilon_{b-2} n^{b-2}$ (assuming EVAL succeeded on $t_{b-2}^{i_{b-1} j_{b-1}}$ in line 2(a)) above and beyond the bounds $t_{b-1}^{ij}(\rho_1, \ldots, \rho_n) \in [l_{ij}, u_{ij}]$ we established above.

We conclude that if, for *all* b, i, and j, EVAL succeeds in producing estimates $e_{i_b}^{ij}$, then $\{\rho_1, \ldots, \rho_n\}$ is a feasible solution for P_2, as desired. The probability of this happening is, by the proof of Lem. 13, at least $1 - d^{2k} n^{k-f}$, since EVAL recursively estimates precisely the same terms during its execution[10]. ☐

Proof of Lem. 16. We begin by proving that for any recursive call to LINEARIZE on t_b with valid upper and lower bounds U and L (i.e. $U, L \neq \infty$), respectively, we have for *any* feasible solution $(\rho_1, \ldots, \rho_n)$ to P_2 that

$$t_b(\rho_1, \ldots, \rho_n) \in [L, U] \pm d(d + \sqrt{2}) \left[\sum_{m=1}^{b-1} (\sqrt{2} d)^{b-1-m} \epsilon_m \right] n^b. \tag{17}$$

We prove this by induction on b, following [AKK99]. For base case $b = 1$, the claim is trivial by line 1(b) of the algorithm. Now, assume by induction hypothesis that

$$t_{b-1}^{ij}(\rho_1, \ldots, \rho_n) \in [l_{ij}, u_{ij}] \pm d(d + \sqrt{2}) \left[\sum_{m=1}^{b-2} (\sqrt{2} d)^{b-2-m} \epsilon_m \right] n^{b-1}.$$

By substituting the values of l_{ij} and u_{ij} from line 2(c), we have

$$t_{b-1}^{ij}(\rho_1, \ldots, \rho_n) \in e_{ij} \pm \left(d(d + \sqrt{2}) \left[\sum_{m=1}^{b-2} (\sqrt{2} d)^{b-2-m} \epsilon_m \right] + \epsilon_{b-1} \right) n^{b-1}.$$

We conclude that

$$\begin{aligned}
t_b(\rho_1, \ldots, \rho_n) &= \sum_{ij} \mathrm{Tr}(\sigma_j \rho_i) t_{b-1}^{ij}(\rho_1, \ldots, \rho_n) \\
&\subseteq \left[\sum_{ij} \mathrm{Tr}(\sigma_j \rho_i) e_{ij} \right] + \left(d(d + \sqrt{2}) \left[\sum_{m=1}^{b-2} (\sqrt{2} d)^{b-2-m} \epsilon_m \right] + \epsilon_{b-1} \right) \left[\sum_{ij} \mathrm{Tr}(\sigma_j \rho_i) \right] n^{b-1} \\
&\subseteq \left[\sum_{ij} \mathrm{Tr}(\sigma_j \rho_i) e_{ij} \right] + \sqrt{2} d \left(d(d + \sqrt{2}) \left[\sum_{m=1}^{b-2} (\sqrt{2} d)^{b-2-m} \epsilon_m \right] + \epsilon_{b-1} \right) n^b \qquad (18) \\
&\subseteq \left[[L, U] \pm \epsilon_{b-1} d^2 n^b \right] + \sqrt{2} d \left(d(d + \sqrt{2}) \left[\sum_{m=1}^{b-2} (\sqrt{2} d)^{b-2-m} \epsilon_m \right] + \epsilon_{b-1} \right) n^b \\
&\subseteq [L, U] \pm d(d + \sqrt{2}) \left[\sum_{m=1}^{b-1} (\sqrt{2} d)^{b-1-m} \epsilon_m \right] n^b,
\end{aligned}$$

where the third statement follows from a calculation similar to Eqn. (15), and the fourth statement from line 3(b) of Alg. 14. This proves the claim of Eqn. (17).

---

[10] This holds even though on line 1 of Alg. 12, we only estimate $d^2 |S|$ of the terms e_{ij} (i.e. EVAL does not actually estimate *all* terms in the recursive decomposition of t_k, as it does not need to) — this is because in our analysis of the probability bound for Alg. 12, we actually produced a looser bound by assuming all n terms e_{ij} are estimated.

To complete the proof of Lem. 16, observe that by Lem. 15, the assignment ρ^{opt} is feasible for P_2 with probability at least $1 - d^{2k}n^{k-f}$. Thus, plugging ρ^{opt} into each of the d^2n linear constraints produced by the recursive calls to LINEARIZE on each t^{ij}_{k-1}, we have by Eqns. (17) and (18) that (with probability $1 - d^{2k}n^{k-f}$)

$$\begin{aligned}
\text{OPT}_P = t_k(\rho^{\text{opt}}) &= \sum_{ij} \text{Tr}\left(\sigma_j \rho_i^{\text{opt}}\right) t^{ij}_{k-1}(\rho^{\text{opt}}) \\
&\subseteq \left[\sum_{i,j} \text{Tr}(\sigma_j \rho_i^{\text{opt}}) e_{ij}\right] \pm \sqrt{2}d \left(d(d+\sqrt{2}) \left[\sum_{m=1}^{k-2} (\sqrt{2}d)^{k-2-m} \epsilon_m\right] + \epsilon_{k-1}\right) n^k \\
&\subseteq \text{OPT}_2 \pm d(d+\sqrt{2}) \left[\sum_{m=1}^{k-1} (\sqrt{2}d)^{k-1-m} \epsilon_m\right] n^k,
\end{aligned}$$

where the last statement follows since ρ^{opt} is not necessarily the optimal solution to P_2. ☐